

AMENDMENTS TO CLAIMS

1. (previously presented) A storage system comprising a redundant array of peer-addressable, multicast storage areas accessible as peers by a first client having a first RAID controller and accessible as peers by a second client having a second RAID controller independent from the first RAID controller.
2. (previously presented) The storage system of claim 1, wherein:
the multicast storage areas are adapted to communicate across a network via split-ID packets comprising both an encapsulating packet and an encapsulated packet; and
each split-ID packets also includes an identifier that is split such that a portion of the identifier is obtained from the encapsulated packet while another portion is obtained from a header portion of the encapsulating packet.
3. (original) The storage system of claim 1, wherein the storage areas of the redundant array share a common multicast address.
4. (original) The storage system of claim 1, comprising a plurality of RAID sets wherein each raid set comprises a plurality of storage areas sharing a common multicast address.
5. (previously presented) A network comprising a first device, a second device, and a plurality of storage devices wherein the first device stores a unit of data on a peer-addressable storage area located on each of the storage devices via a single multicast packet; wherein the second device stores a second unit of data on each of the storage devices via a second single multicast packet; and wherein first device and the second device do not require a common RAID controller to store data on the plurality of storage devices.
6. (previously presented) A network of multicast devices which are operable with a plurality of clients each having an independent RAID controller, and which disaggregate at least one RAID function across multiple peer-addressable, multicast addressable storage areas.
7. (original) The network of claim 6 wherein the at least one RAID function is also disaggregated across multiple device controllers.

8. (previously presented) A storage system comprising a redundant array of peer-addressable, multicast storage areas accessible as peers by a plurality of clients having independent RAID controllers, wherein the system supports auto-annihilation of mooted read requests by disregarding such requests.

9. (canceled)

10. (previously presented) The system of claim 8 wherein auto-annihilation comprises a device that received a read request disregarding the read request if a response to the read request from another device is detected.

11. (previously presented) A storage system comprising a dynamic mirror accessible by a first client having a first RAID controller and accessible by a second client having a second RAID controller independent from the first RAID controller, and wherein the dynamic mirror comprises peer-addressable storage areas that are accessible as peers by the first and the second client.

12. (original) The storage system of claim 11 wherein the dynamic mirror includes a mirrored storage area and at least one corresponding map of incomplete writes.

13. (original) The storage system of claim 11 wherein the dynamic mirror comprises N storage devices and M maps of incomplete writes where M is at least 1 and at most $2*N$.

14. (original) The storage system of claim 13 wherein the map comprises a set of entries wherein each entry is either an LBA or a hash of an LBA of a storage block of a storage area being mirrored.

15. (original) The system of claim 13 comprising at least one process monitoring storage area ACKs sent in response to write commands, the process updating any map associated with a particular area whenever a write command applicable to the area is issued, the process also sending an ACK on behalf of any storage area for which the process did not detect an ACK.

16. (previously presented) The system of claim 15 wherein updating a map comprises setting a flag whenever an ACK is not received and clearing a flag whenever an ACK is received.